

A Directional Diffusion Algorithm for Inpainting

Jan Deriu, Rolf Jagerman, Kai-En Tsay

Department of Computer Science, ETH Zürich, Switzerland

Abstract—The problem of inpainting involves reconstructing the missing areas of an image. Inpainting has many applications, such as reconstructing old damaged photographs or removing obfuscations from images. In this paper we present the directional diffusion algorithm for inpainting. Typical diffusion algorithms are bad at propagating edges from the image into the unknown masked regions. The directional diffusion algorithm improves on the regular diffusion algorithm by reconstructing edges more accurately. It scores better than regular diffusion when reconstructing images that are obfuscated by a text mask.

I. INTRODUCTION

Automatically reconstructing the missing areas of an image is an inference task often referred to as inpainting. It is commonly used to restore old damaged photographs or to remove objects from an image [1]. A variety of approaches have been proposed in the literature to solve this problem. Many successful algorithms use a diffusion-based approach where local pixel information is propagated into the missing areas of an image [2]. Other approaches involve the reconstruction of the full image based on wavelets or dictionaries [3], [4].

In this paper we introduce a novel algorithm that solves the inpainting task by using directional diffusion. This approach is based on the fast diffusion algorithm described by McKenna et al. [5]. It attempts to improve on the fast diffusion algorithm by taking into account the directionality of parts of the image to select proper diffusion kernels.

The paper is structured in several sections. First we will describe the methods we use to accomplish the inpainting task in section II. Next we will discuss the results of our methods in section III and compare them to several baselines. In section IV the strengths and weaknesses of our novel approach are discussed. Finally the paper is concluded in section V.

II. METHODS

We present two algorithms that accomplish the inpainting task: regular diffusion and directional diffusion. The first is a naive yet fast approach. The second is an extension of the first, which takes into account the directionality of image patches to perform smarter diffusions. Although it is slower, it typically leads to better results.

A. Regular diffusion

The idea of diffusion is to fix the known regions of the image and let them diffuse into the unknown regions. This

$$K_{\text{diamond}} = \begin{bmatrix} 0 & 0.25 & 0 \\ 0.25 & 0 & 0.25 \\ 0 & 0.25 & 0 \end{bmatrix}$$

$$K_{\text{diag}} = \begin{bmatrix} 0.38 & 0.04 & 0.04 \\ 0.04 & 0 & 0.04 \\ 0.04 & 0.04 & 0.38 \end{bmatrix}$$

Figure 1: Kernels used for diffusion.

can be done very efficiently by using a convolutional kernel. By iteratively convolving a kernel with the entire image and then restoring the known pixels we can obtain an inpainted image. This process is described in algorithm 1. The quality of the solution heavily depends on the kernel used. A variety of kernels can be used, two of which are displayed in figure 1. To convolve a kernel with an image we have to refer to pixels outside of the border of the image. We replicate the borders of the image outwards so the kernel can properly refer to these values.

Input: Image I , mask M , kernel K and threshold ϵ

Result: Reconstructed image I_r

$K \leftarrow \frac{K}{\sum_i \sum_j K_{i,j}}$ (normalize K to preserve energy);

$I_{\text{prev}} \leftarrow 0_{\text{size}(I)}$;

$I_r \leftarrow I$;

while $\|I_r - I_{\text{prev}}\|_F > \epsilon$ **do**

$I_{\text{prev}} \leftarrow I_r$;

$I_r \leftarrow \text{convolve}(I_r, K)$;

$I_r \leftarrow I_r \circ \mathbf{1}_{M=0} + I \circ \mathbf{1}_{M \neq 0}$;

end

Algorithm 1: Diffusion algorithm for inpainting. We denote element-wise multiplication with the \circ operator. The $\mathbf{1}_{M=0}$ function represents a matrix with elements (i, j) set to 1 when $M_{i,j} = 0$ and 0 otherwise.

B. Directional diffusion

Regular diffusion will present noticeable artifacts. These typically occur at high-contrast edges of the image. Because the diffusion kernel equally weighs pixels from all directions it creates a blur effect and is unable to properly propagate edges into the unknown regions. This effect can be seen in figure 2a. To help resolve this problem we use directional kernels that weigh pixels from certain directions more

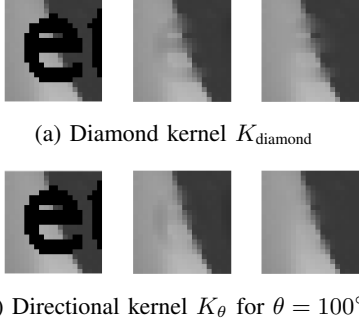


Figure 2: Step-by-step illustration of the diffusion process with different kernels. Each step represent 20 iterations.

heavily. We define directionality to be the angle in which the high-contrast edges of an image are directed. An image with mostly horizontal lines will have a directionality of 0 degrees whereas an image with mostly vertical lines will have a directionality of 90 degrees. By inferring the correct directionality, the artifacts can be noticeably reduced as seen in figure 2b.

To do this we first apply the regular diffusion algorithm to the image to get an estimate of the original image. Then we divide the image into separate patches. For each patch we infer the directionality using a heuristic. Based on this directionality we construct a directional kernel. Finally we apply this kernel to do the inpainting for that specific patch. We will now explain these steps in more detail.

1) *Divide image into $n \times n$ patches:* The directionality of the contents of an image is typically a local property. Hence we want to know what the directionality of a small image patch is. An image I can be broken down into patches of size $n \times n$. We define a patch P starting at pixel location i, j as:

$$P = \begin{bmatrix} I_{i,j} & I_{i+1,j} & \dots & I_{i+n,j} \\ I_{i,j+1} & I_{i+1,j+1} & \dots & I_{i+n,j+1} \\ \vdots & \vdots & \ddots & \vdots \\ I_{i,j+n} & I_{i+1,j+n} & \dots & I_{i+n,j+n} \end{bmatrix}$$

2) *Infer directionality of an image patch:* Given an image patch P we wish to compute an estimated angle θ of the contents of this patch. Intuitively this means for patches with lots of horizontal stripes we would get $\theta \approx 0^\circ$ and for patches with lots of vertical stripes we would get $\theta \approx 90^\circ$. The first thing we need is a shift operator for matrices. Given a matrix $M \in \mathbb{R}^{n \times m}$ we can compute the shifted matrix $M^{(x,y)}$, which circularly shifts columns towards the right x times and rows towards the bottom y times:

$$M_{i,j}^{(x,y)} = M_{(i+x) \bmod n, (j+y) \bmod m}$$

Using this shift operator, we compute two basic metrics on the image patch:

$$v = \sum_{i,j} |P_{i,j} - P_{i,j}^{(1,0)}|$$

$$h = \sum_{i,j} |P_{i,j} - P_{i,j}^{(0,1)}|$$

The first metric v checks how much each pixel differs from the one left of it and provides large values for images with a lot of vertical stripes. The second metric h does the same but checks how much each pixel differs from the one below it. This gives large values for images with a lot of horizontal stripes.

By comparing h and v we can determine if the contents of the image patch will have mostly vertical lines ($v > h$), mostly horizontal lines ($v < h$) or diagonal lines ($v \approx h$). We can turn this into an angle, by computing the fraction of h to $h + v$. We add 1 to both sides to prevent division by zero:

$$\theta_1 = 90 \frac{h + 1}{h + v + 1}$$

This approach only gives an angle from 0° to 90° . It cannot infer if the diagonal lines on the image patch are left-slanted or right-slanted since these will have $h \approx v$. To help resolve this we propose a heuristic for the diagonal slant by computing d , which checks how much each pixel differs from the one diagonally right-below it. This gives large values for images with right-slanted diagonal lines but small values for images with left-slanted diagonal lines. We wish to normalize this value by dividing it by $v + h$, so it ranges from 0 to 1. We add 1 to both sides to prevent division by zero:

$$d = \frac{1 + \sum_{i,j} |P_{i,j} - P_{i,j}^{(1,1)}|}{1 + v + h}$$

If this value d is close to 1, we can assume the image is right-slanted. If it is close to 0, we can assume the image is left-slanted. A value close to 0.5 indicates that the image patch has both left and right slanted lines. We can now compute the angle θ with the additional parameter d using the following heuristic:

$$\theta = -90 + \begin{cases} 90d + \theta_1 & \text{if } d > 0.6 \\ 90 - \theta_1 & \text{if } d \leq 0.6 \end{cases}$$

In figure 3 we show the angle θ , which has been computed on image patches of size 16×16 , as white lines. The heuristic works well on image patches where the contents is simple and consists of mostly straight lines. When the content of an image patch gets more complex, inferring the angle becomes more difficult and the heuristic fails.



(a) The Elaine image.

(b) The splash image.

(c) The aerial image.

Figure 3: The directionality of 16x16 pixel patches shown as white lines. Images obtained from the SIPI database [6]

3) *Constructing a directional kernel*: After computing the angle θ for an image patch, we wish to construct a directional kernel K_θ , which puts more weight on pixels that align with the angle θ . We start with a diagonal kernel K_{diag} (see figure 1). This kernel is converted to a 3×3 image which is rotated by $\theta + 45$ degrees (because the initial kernel was diagonal, we add 45 degrees). It is then cropped to obtain a new 3×3 image representing the rotated kernel. During the rotation process, bicubic interpolation is used to smooth the intermediate values [7]. The resulting image is converted back to a 3×3 matrix representing the kernel K_θ .

4) *Per-patch diffusion using K_θ* : Now that we have a kernel K_θ for each image patch P we can run the diffusion algorithm (see algorithm 1) on each of the patches. The reconstructed image is created by putting all the patches back together after running the diffusion algorithm on them.

III. RESULTS

We test our methods on seventeen different 512×512 grayscale images using six different algorithms:

- 1) Sparse-coding with a DCT dictionary [8].
- 2) Sparse-coding with a Haar wavelet [3].
- 3) Singular Value Decomposition [9].
- 4) Regular diffusion with a K_{diamond} kernel.
- 5) Directional diffusion with patches of size 16×16 .
- 6) Directional diffusion with patches of size 32×32 .

To test the performance of the algorithm, we consider two types of masks:

- 1) Nine different masks of randomly missing pixels distributed uniformly. These range from 10% to 90% missing pixels. This type of mask has many small missing regions.
- 2) A mask represented by a sample of text. This type of mask has many medium-sized missing regions.

We have set up the experiments as follows: We take the original image I and a mask M . We construct a damaged

Algorithm	MSE	Runtime
Directional Diffusion (16×16)	0.00055 \pm 0.00051	8.7 \pm 2.13
Directional Diffusion (32×32)	0.00057 \pm 0.00053	2.4 \pm 0.04
Diffusion (K_{diamond})	0.00061 \pm 0.00057	0.5 \pm 0.06
Sparse-coding (DCT)	0.0015 \pm 0.0012	12.8 \pm 4.67
Sparse-coding (Haar wavelet)	0.0024 \pm 0.0021	13.0 \pm 3.72
Singular Value Decomposition	0.0019 \pm 0.0018	0.7 \pm 0.09

Table I: Mean squared error and runtime (in seconds) across different algorithms for the text mask. The best result is highlighted in bold.

image I_{damaged} based on I but with the pixels indicated by the mask M set to the fixed value 0. We run the relevant inpainting algorithms by providing it with both the damaged image and the mask. The algorithm will return a recovered image I^{rec} , which we can compare to the original image I .

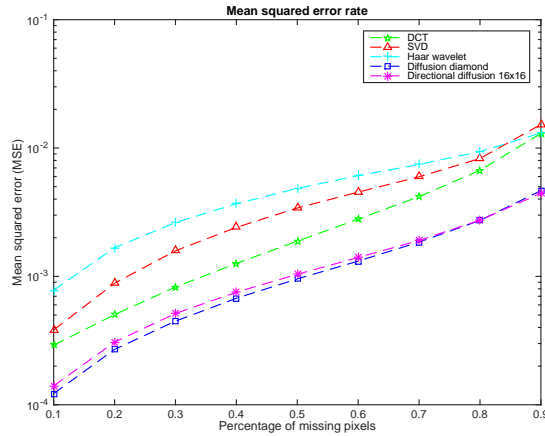
We compare the algorithms based on two criteria: the mean-squared error and the runtime of the algorithm. The mean-squared error is computed as the mean of the square of the difference in all pixel intensity values of the original image and the reconstructed image:

$$\text{MSE}(I, I^{\text{rec}}) = \frac{1}{512 \cdot 512} \sum_{i,j} (I_{i,j} - I_{i,j}^{\text{rec}})^2$$

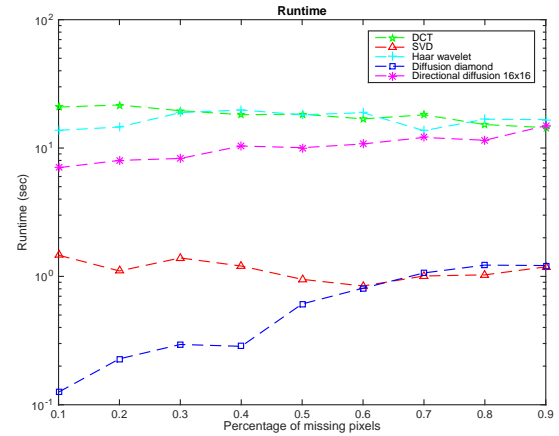
The mean squared error of the algorithms for the randomly generated masks are displayed in figure 4a and the runtime is displayed in figure 4b. The mean squared error and runtime of the algorithms for the text mask is given in table I.

IV. DISCUSSION

Both the regular diffusion and the directional diffusion algorithm show promising results. From figure 4a we see that the regular diffusion algorithm with the K_{diamond} kernel works best on a mask with randomly missing pixels. This can be explained by the fact that this algorithm diffuses



(a) Mean squared error in log scale.



(b) Runtime in log scale.

Figure 4: Mean squared error and runtime comparison of different algorithms. To improve readability we only consider the directional diffusion algorithm for patches of size 16×16 .

nearby pixels into the missing regions. A mask with randomly missing pixels will, on average, have at least some pixels in the direct or near neighborhood of a pixel that we try to inpaint.

Table I shows the mean squared error of the algorithms on a structured mask, namely a piece of text. The regular diffusion algorithm does not perform as well in this setting. This is because it relies on nearby known pixels which are less common in masks with medium-sized areas of missing pixels. As a consequence, the high-contrasting edges of the underlying image are not properly extended into the unknown regions. The directional diffusion algorithm helps resolve this by aligning the kernel K_θ with the directionality of the image patches. Because of this the directional diffusion algorithm scores best.

V. CONCLUSION

In this paper we presented the directional diffusion algorithm for the inpainting problem. It is an extension to the regular diffusion algorithm. Directional diffusion outperforms regular diffusion when applied to text masks.

The main drawback of the directional diffusion algorithm is its runtime. However, due to the nature of the algorithm, it is very easy to parallelize. By applying the patch-dependent kernels K_θ on all patches simultaneously we can achieve great speed-ups.

Additionally, the heuristic used to infer the directionality θ is not perfect and can be improved. One could use the sobel operator to compute gradients of an image patch [10]. Based on these gradients it might be possible to get a more robust estimate of the directionality θ of the image patches.

REFERENCES

- [1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, “Image inpainting,” in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 2000, pp. 417–424.
- [2] L. Alvarez, P.-L. Lions, and J.-M. Morel, “Image selective smoothing and edge detection by nonlinear diffusion. ii,” *SIAM Journal on numerical analysis*, vol. 29, no. 3, pp. 845–866, 1992.
- [3] T. Hofmann, B. McWilliams, and J. Buhmann, “Sparse coding, classical methods,” 2015, <http://cil.inf.ethz.ch/material/lecture/lecture08.pdf>.
- [4] —, “Sparse coding, learned dictionaries,” 2015, <http://cil.inf.ethz.ch/material/lecture/lecture10.pdf>.
- [5] M. M. O. B. B. Richard and M. Y.-S. Chang, “Fast digital image inpainting,” 2001.
- [6] “Sipi image database, volume 3: Miscellaneous,” <http://sipi.usc.edu/database/database.php?volume=misc>, accessed: 2015-11-11.
- [7] R. G. Keys, “Cubic convolution interpolation for digital image processing,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 29, no. 6, pp. 1153–1160, 1981.
- [8] T. Hofmann, B. McWilliams, and J. Buhmann, “Sparse coding, overcomplete dictionaries,” 2015, <http://cil.inf.ethz.ch/material/lecture/lecture09.pdf>.
- [9] T. Hofmann, “Singular value decomposition,” 2015, <http://cil.inf.ethz.ch/material/lecture/lecture03.pdf>.
- [10] I. Sobel, “History and definition of the sobel operator,” Retrieved from <http://www.researchgate.net/>, 2014.